

# Fast Fourier Color Constancy

Jonathan T. Barron

barron@google.com

Yun-Ta Tsai

yuntatsai@google.com

## Abstract

We present *Fast Fourier Color Constancy (FFCC)*, a color constancy algorithm which solves illuminant estimation by reducing it to a spatial localization task on a torus. By operating in the frequency domain, FFCC produces lower error rates than the previous state-of-the-art by 10 – 13% while being 250 – 3000 $\times$  faster. This unconventional approach introduces challenges regarding aliasing, directional statistics, and preconditioning, which we address. By producing a complete posterior distribution over illuminants instead of a single illuminant estimate, FFCC enables better training techniques, an effective temporal smoothing technique, and richer methods for error analysis. Our implementation of FFCC runs at  $\sim 700$  frames per second on a mobile device, allowing it to be used as an accurate, real-time, temporally-coherent automatic white balance algorithm.

## 1. Intro

A fundamental problem in computer vision is that of estimating the underlying world that resulted in some observed image [1, 5]. One subset of this problem is color constancy: estimating the color of the illuminant of the scene and the colors of the objects in the scene viewed under a white light. Despite its apparent simplicity, this problem has yielded a great deal of depth and challenge for both the human vision and computer vision communities [15, 20]. Color constancy is also a practical concern in the camera industry: producing a natural looking photograph without user intervention requires that the illuminant be automatically estimated and discounted, a process referred to as “auto white balance” among practitioners. Though there is a profound historical connection between color constancy and consumer photography (exemplified by Edwin Land, the inventor of both Retinex theory [24] and the Polaroid instant camera), “color constancy” and “white balance” have come to mean different things — color constancy aims to recover the veridical world behind an image, while white balance aims to give an image a pleasant appearance consistent with some aesthetic or cultural norm. But with the current ubiquity of learning-

based techniques in computer vision, both problems reduce to just estimating the “best” illuminant from an image, and the question of whether that illuminant is objectively true or subjectively attractive is just a matter of the data used during training.

Despite their accuracy, modern learning-based color constancy algorithms are not immediately suitable as practical white balance algorithms, as practical white balance has several requirements besides accuracy:

**Speed** - An algorithm running in a camera’s viewfinder must run at 30 FPS on mobile hardware. But a camera’s compute budget is precious: demosaicing, face detection, auto exposure, etc, must also run simultaneously and in real time. Spending more than a small fraction (say, 5 – 10%) of a camera’s compute budget on white balance is impractical, suggesting that our speed requirement is closer to 1 – 5 milliseconds per frame.

**Impoverished Input** - Most color constancy algorithms are designed for full resolution, high bit-depth input images, but operating on such large images is challenging and costly in practice. To be fast, the algorithm must work well on the small, low bit-depth “preview” images ( $32 \times 24$  or  $64 \times 48$  pixels, 8-bit) which are usually computed by specialized camera hardware for this task.

**Uncertainty** - In addition to the illuminant, the algorithm should produce some confidence measure or a complete posterior distribution over illuminants, thereby enabling convenient downstream integration with hand-engineered heuristics or external sources of information.

**Temporal Coherence** - The algorithm should allow the estimated illuminant to be smoothed over time, to prevent color composition in videos from varying erratically.

In this paper we present a novel color constancy algorithm, which we call “Fast Fourier Color Constancy” (FFCC). Viewed as a color constancy algorithm, FFCC is 10 – 13% more accurate than the state of the art on standard benchmarks. Viewed as a prospective white balance algorithm, FFCC addresses our previously described requirements: Our technique is 250–3000 $\times$  faster than the state of the art, and is capable of running at 1.44 milliseconds per frame on a standard consumer mobile platform using the thumbnail images already produced by that camera’s hard-

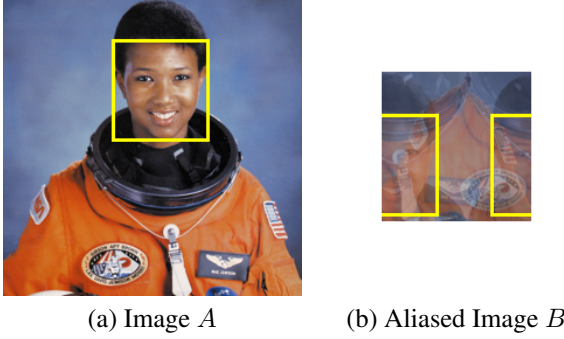


Figure 1: CCC [4] reduces color constancy to a 2D localization problem similar to object detection (1a). FFCC repeatedly wraps this 2D localization problem around a small torus (1b), which creates challenges but allows for faster illuminant estimation. See the text for details.

ware. FFCC produces a complete posterior distribution over illuminants which allows us to reason about uncertainty and enables simple and effective temporal smoothing.

We build on the “Convolutional Color Constancy” (CCC) approach of [4], which is currently one of the top-performing techniques on standard color constancy benchmarks [11, 18, 28]. CCC works by observing that applying a per-channel gain to a linear RGB image is equivalent to inducing a 2D translation of the log-chroma histogram of that image, which allows color constancy to be reduced to the task of localizing a signature in log-chroma histogram space. This reduction is at the core of the success of CCC and, by extension, our FFCC technique; see [4] for a thorough explanation. The primary difference between FFCC is that instead of performing an expensive localization on a large log-chroma plane, we perform a cheap localization on a small log-chroma *torus*.

At a high level, CCC *reduces* color constancy to object detection — in the computability theory sense of “reduce”. FFCC reduces color constancy to localization on a torus instead of a plane, and because this task has no intuitive analogue in computer vision we will attempt to provide one<sup>1</sup>. Given a large image  $A$  on which we would like to perform object detection, imagine constructing a smaller  $n \times n$  image  $B$  in which each pixel in  $B$  is the sum of all values in  $A$  separated by a multiple of  $n$  pixels in either dimension:

$$B(i, j) = \sum_{l, m} A(i + nl, j + nm) \quad (1)$$

This amounts to taking  $A$  and repeatedly wrapping it around a small torus (see Figure 1). Detecting objects this way may yield a speedup as the image being processed is smaller,

<sup>1</sup> We cannot speak to the merit of this idea in the context of object detection, and we present it here solely to provide an intuition of our work on color constancy

but it also raises new problems: pixel values are corrupted with superimposed shapes that make detection difficult, detections must “wrap” around the edges of this toroidal image, and instead of an absolute, global location we can only recover an aliased, incomplete location. FFCC works by taking the large convolutional problem of CCC (ie, face detection on  $A$ ) and aliasing that problem down to a smaller size where it can be solved efficiently (ie, face detection on  $B$ ). We will show that we can learn an effective color constancy model in the face of the difficulty and ambiguity introduced by aliasing. This convolutional classifier will be implemented and learned using FFTs, because the naturally periodic nature of FFT convolutions resolves the problem of detections “wrapping” around the edge of toroidal images, and produces a significant speedup.

Our approach to color constancy introduces a number of issues. The aforementioned periodic ambiguity resulting from operating on a torus (which we dub “illuminant aliasing”) requires novel new techniques for recovering a global illuminant estimate from an aliased estimate (Section 3). Localizing the centroid of the illuminant on a torus is difficult, requiring that we adopt and extend techniques from the directional statistics literature (Section 4). But our approach presents a number of benefits. Our model improves accuracy relative to CCC by 10 – 13% while retaining its flexibility, and allows us to construct priors over illuminants (Section 5). By learning in the frequency-domain we can construct a novel method for fast frequency-domain regularization and preconditioning, making FFCC training  $15\times$  faster than CCC (Section 6). Our model produces a complete posterior over illuminants as output, allowing us to construct a Kalman filter-like approach for processing videos instead of independent images (Section 7).

## 2. Convolutional Color Constancy

Let us review the assumptions made in CCC and inherited by our model. Assume that we have a photometrically linear input image  $I$  from a camera, with a black level of zero and with no saturated pixels<sup>2</sup>. Each pixel  $k$ ’s RGB value in image  $I$  is assumed to be the product of that pixel’s “true” white-balanced RGB value  $W^{(k)}$  and some global RGB illumination  $L$  shared by all pixels:

$$\forall_k \begin{bmatrix} I_r^{(k)} \\ I_g^{(k)} \\ I_b^{(k)} \end{bmatrix} = \begin{bmatrix} W_r^{(k)} \\ W_g^{(k)} \\ W_b^{(k)} \end{bmatrix} \circ \begin{bmatrix} L_r \\ L_g \\ L_b \end{bmatrix} \quad (2)$$

The task of color constancy is to use the input image  $I$  to estimate  $L$ , and with that produce  $W^{(k)} = I^{(k)} / L$ .

Given a pixel from our input RGB image  $I^{(k)}$ , CCC de-

<sup>2</sup>in practice, saturated pixels are identified and removed from all downstream computation, similarly to how color checker pixels are ignored.

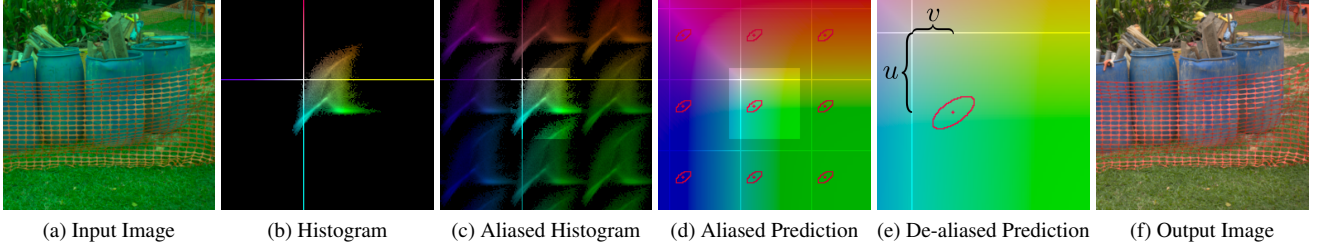


Figure 2: An overview of our pipeline demonstrating the problem of illuminant aliasing. Similarly to CCC, we take an input image (2a) and transform it into a log-chroma histogram (2b shown here in the same format as [4]). But unlike CCC, our histograms are small and toroidal, meaning that pixels can “wrap around” the edges (2c, with the torus “unwrapped” once in every direction). This means that the centroid of a filtered histogram, which would simply be the illuminant estimate in CCC, is instead an infinite family of possible illuminants (2d). This requires *de-aliasing*, some technique for disambiguating between illuminants to select the single most likely estimate (2e, shown as a point surrounded by an ellipse visualizing the output covariance of our model). Our model’s output  $(u, v)$  coordinates in this de-aliased log-chroma space corresponds to the color of the illuminant, which can then be divided into the input image to produce a white balanced image (2f).

finds two log-chroma measures:

$$u^{(k)} = \log \left( I_g^{(k)} / I_r^{(k)} \right) \quad v^{(k)} = \log \left( I_g^{(k)} / I_b^{(k)} \right) \quad (3)$$

The absolute scale of  $L$  is assumed to be unrecoverable, so estimating  $L$  simply requires estimating its log-chroma:

$$L_u = \log (L_g / L_r) \quad L_v = \log (L_g / L_b) \quad (4)$$

After recovering  $(L_u, L_v)$ , assuming that  $L$  has a magnitude of 1 lets us recover the RGB values of the illuminant:

$$L_r = \frac{\exp(-L_u)}{z} \quad L_g = \frac{1}{z} \quad L_b = \frac{\exp(-L_v)}{z} \quad (5)$$

$$z = \sqrt{\exp(-L_u)^2 + \exp(-L_v)^2 + 1}$$

Framing color constancy in terms of predicting log-chroma has several small advantages over the standard RGB approach (2 unknowns instead of 3, better numerical stability, etc) but the primary advantage of this approach is that using log-chroma turns the multiplicative constraint relating  $W$  and  $I$  into an additive constraint, and this in turn enables a convolutional approach to color constancy. As shown in [4], color constancy can be framed as a 2D spatial localization task on a log-chroma histogram  $N$ , where some sliding-window classifier is used to filter that histogram and centroid of that filtered histogram is used as the log-chroma of the illuminant.

### 3. Illuminant Aliasing

We assume the same convolutional premise of CCC, but with one primary difference to improve quality and speed: we use FFTs to perform the convolution which filters the log-chroma histogram, and we use a small histogram to make that convolution as fast as possible. This change may seem trivial, but the periodic nature of FFT convolution

combined with the properties of natural images has a significant effect, as we will demonstrate.

Similarly to CCC, given an input image  $I$  we construct a histogram  $N$  from  $I$ , where  $N(i, j)$  is the number of pixels in  $I$  whose log-chroma is near the  $(u, v)$  coordinates corresponding to histogram position  $(i, j)$ :

$$N(i, j) = \sum_k \left( \text{mod} \left( \frac{u^{(k)} - u_{lo}}{h} - i, n \right) < 1 \right. \\ \left. \wedge \text{mod} \left( \frac{v^{(k)} - v_{lo}}{h} - j, n \right) < 1 \right) \quad (6)$$

Where  $i, j$  are 0-indexed,  $n = 64$  is the number of bins,  $h = 1/32$  is the bin size, and  $(u_{lo}, v_{lo})$  is the starting point of the histogram. Because our histogram is too small to contain the wide spread of colors present in most natural images, we use modular arithmetic to cause pixels to “wrap around” with respect to log-chroma (any other standard boundary condition would violate our convolutional assumption and would cause many image pixels to be ignored). This means that, unlike standard CCC, a single  $(i, j)$  coordinate in the histogram no longer corresponds to an absolute  $(u, v)$  color, but instead corresponds to an infinite family of  $(u, v)$  colors. Accordingly, the centroid of a filtered histogram no longer corresponds to the color of the illuminant, but instead is an infinite set of illuminants. We will refer to this phenomenon as *illuminant aliasing*. Solving this problem requires that we use some technique to de-alias an aliased illuminant estimate<sup>3</sup>. A high-level

<sup>3</sup>It is tempting to refer to resolving the illuminant aliasing problem as “anti-aliasing”, but anti-aliasing usually refers to preprocessing a signal to prevent aliasing during some resampling operation, which does not appear possible in our framework. “De-aliasing” suggests that we allow aliasing to happen to the input, but then remove the aliasing from the output.

outline of our FFCC pipeline which illustrates illuminant (de-)aliasing can be seen in Fig. 2.

De-aliasing requires that we use some external information (or some external color constancy algorithm) to disambiguate between illuminants. An intuitive approach is to select the illuminant which causes the average image color to be as neutral as possible, which we call “gray world de-aliasing”. We compute average log-chroma values  $(\bar{u}, \bar{v})$  for the entire image and use this to turn an aliased illuminant estimate  $(\hat{L}_u, \hat{L}_v)$  into a de-aliased illuminant  $(\hat{L}'_u, \hat{L}'_v)$ :

$$\bar{u} = \log(\sum_k u^{(k)}) \quad \bar{v} = \log(\sum_k v^{(k)}) \quad (7)$$

$$\begin{bmatrix} \hat{L}'_u \\ \hat{L}'_v \end{bmatrix} = \begin{bmatrix} \hat{L}_u \\ \hat{L}_v \end{bmatrix} - (nh) \left[ \frac{1}{nh} \begin{bmatrix} \hat{L}_u - \bar{u} \\ \hat{L}_v - \bar{v} \end{bmatrix} + \frac{1}{2} \right] \quad (8)$$

Another approach, which we call “gray light de-aliasing”, is to assume that the illuminant is as close to the center of the histogram as possible. This de-aliasing approach simply requires carefully setting the starting point of the histogram  $u_{lo} = v_{lo}$  such that the true illuminants in natural scenes all lie within the span of the histogram, and setting  $\hat{L}' = \hat{L}$ . We do this by setting  $u_{lo} = v_{lo}$  to maximize the distance between the edges of the histogram and the bounding box surrounding the ground-truth illuminants in the training data<sup>4</sup>. Gray light de-aliasing is trivial to implement but, unlike gray world de-aliasing, it will systematically fail if the histogram is too small to fit all illuminants within its span.

To summarize the difference between CCC [4] and our approach with regards to illuminant aliasing, CCC (approximately) performs illuminant estimation as follows:

$$\begin{bmatrix} \hat{L}_u \\ \hat{L}_v \end{bmatrix} = \begin{bmatrix} u_{lo} \\ v_{lo} \end{bmatrix} + h \left( \arg \max_{i,j} (N * F) \right) \quad (9)$$

Where  $N * F$  is performed using a pyramid convolution. FFCC corresponds to this procedure:

$$P \leftarrow \text{softmax}(N * F) \quad (10)$$

$$(\mu, \Sigma) \leftarrow \text{fit\_bvm}(P) \quad (11)$$

$$\begin{bmatrix} \hat{L}_u \\ \hat{L}_v \end{bmatrix} \leftarrow \text{de\_alias}(\mu) \quad (12)$$

Where  $N$  is a small, aliased, toroidal histogram, convolution is performed with FFTs, and the centroid of the filtered histogram is carefully extracted and de-aliased as necessary. Because this pipeline is differentiable, we can train our model in an end-to-end fashion by propagating the gradients of some loss computed on the de-aliased illuminant

<sup>4</sup>Our histograms are shifted toward green colors rather than centered around a neutral color, as cameras are traditionally designed with an more sensitive green channel which enables white balance to be performed by gaining red and blue up without causing color clipping. Ignoring this practical issue, our approach can be thought of as centering our histograms around a neutral white light

prediction  $\hat{L}$  back onto the learned filters  $F$ . The centroid fitting in Eq. 11 is performed by fitting a bivariate von Mises distribution to a PDF, which we will now explain.

#### 4. Differentiable Bivariate von Mises

Our architecture requires some mechanism for reducing a per-color toroidal PDF  $P(i, j)$  to a single estimate of the illuminant. Localizing the center of mass of a histogram defined on a torus is difficult: fitting a bivariate Gaussian may fail when the input distribution “wraps around” the sides of the PDF, as shown in Fig. 3. Additionally, for the sake of temporal smoothing (Section 7) and confidence estimation, we want our model to predict a well-calibrated covariance matrix around the center of mass of  $P$ . This requires that our model be trained end-to-end, which therefore requires that our mean/covariance fitting be analytically differentiable and therefore usable as a “layer” in our deep learning architecture. To address these problems we present a variant of the bivariate von Mises distribution [25], which we will use to efficiently localize the mean and covariance of  $P$  in a manner that allows for easy backpropagation.

The bivariate von Mises distribution (BVM) is a common parameterization of a PDF on a torus. There exist several parametrizations which mostly differ in how “concentration” is represented (“concentration” having a similar meaning to covariance). All of these parametrizations present problems in our use case: none have closed form expressions for maximum likelihood estimators [22], none lend themselves to convenient backpropagation, and all define concentration in terms of angles and therefore require “conversion” to covariance matrices during color de-aliasing. For these reasons we present an alternative parametrization in which we directly estimate a BVM as a mean  $\mu$  and covariance  $\Sigma$  in a simple and differentiable closed form expression. Though necessarily approximate, our estimator is accurate when the distribution is well-concentrated, which is generally the case for our task.

Our input is a PDF  $P(i, j)$  of size  $n \times n$ , where  $i$  and  $j$  are integers in  $[0, n - 1]$ . For convenience we define a mapping from  $i$  or  $j$  to angles in  $[0, 2\pi)$  and the marginal distributions of  $P$  with respect to  $i$  and  $j$ :

$$\theta(i) = \frac{2\pi i}{n} \quad P_i(i) = \sum_j P(i, j) \quad P_j(j) = \sum_i P(i, j)$$

We also define the marginal expectation of the sine and cosine of the angle:

$$y_i = \sum_i P_i(i) \sin(\theta(i)) \quad x_i = \sum_i P_i(i) \cos(\theta(i)) \quad (13)$$

With  $x_j$  and  $y_j$  defined similarly.

Estimating the mean  $\mu$  of a BVM from a histogram just



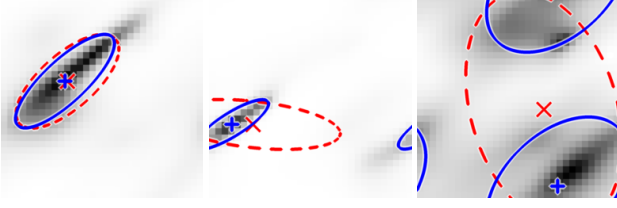


Figure 3: We fit a bivariate von Mises distribution (shown in solid blue) to toroidal PDFs  $P(i, j)$  to produce an aliased illuminant estimate. Contrast this with fitting a bivariate Gaussian (shown in dashed red) which treats the PDF as if it lies on a plane. Both approaches behave similarly if the distribution lies near the center of the unwrapped plane (left) but fitting a Gaussian fails as the distribution begins to “wrap around” the edge (middle, right).

requires computing the circular mean in  $i$  and  $j$ :

$$\boldsymbol{\mu} = \begin{bmatrix} u_{lo} \\ v_{lo} \end{bmatrix} + h \begin{bmatrix} \text{mod} \left( \frac{n}{2\pi} \text{atan2}(y_i, x_i), n \right) \\ \text{mod} \left( \frac{n}{2\pi} \text{atan2}(y_j, x_j), n \right) \end{bmatrix} \quad (14)$$

Eq. 14 includes gray light de-aliasing, though gray world de-aliasing can also be applied to  $\boldsymbol{\mu}$  after fitting.

We can fit the covariance of our model by simply “unwrapping” the coordinates of the histogram relative to the estimated mean and treating these unwrapped coordinates as though we are fitting a bivariate Gaussian. We define the “unwrapped”  $(i, j)$  coordinates such that the “wrap around” point on the torus lies as far away from the mean as possible, or equivalently, such that the unwrapped coordinates are as close to the mean as possible:

$$\begin{aligned} \bar{i} &= \text{mod} \left( i - \left\lfloor \frac{\boldsymbol{\mu}_u - u_{lo}}{h} \right\rfloor + \frac{n}{2}, n \right) \\ \bar{j} &= \text{mod} \left( j - \left\lfloor \frac{\boldsymbol{\mu}_v - v_{lo}}{h} \right\rfloor + \frac{n}{2}, n \right) \end{aligned} \quad (15)$$

Our estimated covariance matrix is simply the sample covariance of  $P(\bar{i}, \bar{j})$ :

$$\mathbb{E}[\bar{i}] = \sum_i P_i(i) \bar{i} \quad \mathbb{E}[\bar{j}] = \sum_j P_j(j) \bar{j} \quad (16)$$

$$\boldsymbol{\Sigma} = h^2 \begin{bmatrix} \epsilon + \sum_i P_i(i) \bar{i}^2 - \mathbb{E}[\bar{i}]^2 & \sum_{i,j} P(i, j) \bar{i} \bar{j} - \mathbb{E}[\bar{i}] \mathbb{E}[\bar{j}] \\ \sum_{i,j} P(i, j) \bar{i} \bar{j} - \mathbb{E}[\bar{i}] \mathbb{E}[\bar{j}] & \epsilon + \sum_j P_j(j) \bar{j}^2 - \mathbb{E}[\bar{j}]^2 \end{bmatrix} \quad (17)$$

We regularize the sample covariance matrix slightly by adding a constant  $\epsilon = 1$  to the diagonal.

With our estimated mean and covariance we can now compute the loss that we minimize during training, which the negative log-likelihood of a Gaussian (ignoring scale factors and constants) relative to the true illuminant  $L^*$ :

$$f(\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \log |\boldsymbol{\Sigma}| + \left( \begin{bmatrix} L_u^* \\ L_v^* \end{bmatrix} - \boldsymbol{\mu} \right)^T \boldsymbol{\Sigma}^{-1} \left( \begin{bmatrix} L_u^* \\ L_v^* \end{bmatrix} - \boldsymbol{\mu} \right) \quad (18)$$

Using this loss causes our model to produce a well-calibrated complete posterior of the illuminant instead of just a single estimate. This posterior will be useful when processing video sequences (Section 7) and also allows us to attach confidence estimates to our predictions using the entropy of  $\boldsymbol{\Sigma}$  (see the supplement).

Our entire system is trained end-to-end in a “deep” fashion, which requires that every step in BVM fitting and loss computation be analytically differentiable. See the supplement for the analytical gradients for Eqs. 14, 17, and 18, which can be chained together to backpropagate the gradient of  $f(\cdot)$  onto the input PDF  $P$ .

## 5. Model Extensions

The system we have described thus far (compute a periodic histogram of each pixel’s log-chroma, apply a learned FFT convolution, apply a softmax, fit a de-aliased bivariate von Mises distribution) works reasonably well (model A in Table 1) but does not produce state-of-the-art results. This is likely because this model reasons about pixels independently, ignores all spatial information in the image, and does not consider the absolute color of the illuminant. Here we present extensions to the model which address these issues and improve accuracy accordingly.

As explored in [4], a CCC-like model can be generalized to a set of “augmented” images provided that these images are non-negative and “scale with intensity” [13]. This allows us to apply certain filtering operations to our image  $I$  and, instead of constructing a single histogram from our image, construct a “stack” of histograms constructed from the image and its filtered versions. Instead of simply learning and applying one filter, we learn a stack of filters and sum across channels after convolution. The general family of augmented images used in [4] are expensive to compute, so we instead use just the input image  $I$  and a local measure of absolute deviation in the input image:

$$E(x, y, c) = \frac{1}{8} \sum_{i=-1}^1 \sum_{j=-1}^1 |I(x, y, c) - I(x+i, y+j, c)| \quad (19)$$

These two features appears to perform similarly to the four features used in [4], while being quicker to compute.

Just as a sliding-window object detector is often invariant to the absolute location of an object in an image, the convolutional nature of our baseline model makes it invariant to any global shift of the color of the input image. This means that our baseline model *cannot* rely on any statistical regularities of the illumination by, say, modeling black body radiation, the specific properties of commonly manufactured light bulbs, or any varying spectral sensitivity across cameras. Though CCC does not model illumination directly, it appears to indirectly reason about illumination by using the boundary conditions of its pyramid convolu-

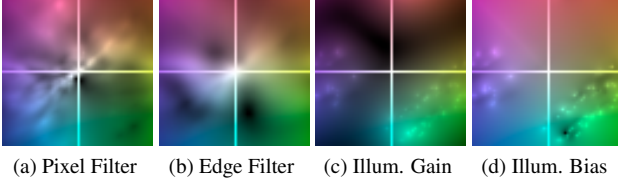


Figure 4: A complete learned model (Model J in Table 1) shown in centered  $(u, v)$  log-chroma space, with brightness indicating larger values. Our learned filters are centered around the origin (the predicted white point) and our illuminant gain and bias maps model the black body curve and varying camera sensitivity as two wrap-around line segments (this dataset consists of images from two different cameras).

tion to learn a model which is not truly spatially varying and is therefore sensitive to absolute color. Because a torus has no boundaries, our model is invariant to global input color so we must introduce a mechanism for directly reasoning about illuminants. We use a per-illuminant “gain” map  $G(i, j)$  and “bias” map  $B(i, j)$ , which together apply a per-illuminant affine transformation to the output of our previously-described convolution at (aliased) color  $(i, j)$ . The bias  $B$  causes our model to prefer certain illuminants over others, while the gain  $G$  causes the contribution of the convolution at certain colors to be amplified, which can be thought of as a kind of informed bias.

Our two extensions (an augmented edge channel and an illuminant gain/bias map) let us redefine the  $P$  in Eq. 10 as

$$P = \text{softmax} \left( B + G \circ \sum_c (N_c * F_c) \right) \quad (20)$$

Where  $\{F_c\}$  are the set of learned filters for each augmented channel’s histogram  $N_c$ ,  $G$  is our learned gain map, and  $B$  is our learned bias map. In practice we actually parametrize  $G_{\log}$  when training and define  $G = \exp(G_{\log})$ , which constraints  $G$  to be non-negative. Visualizations of  $G$  and  $B$  and our learned filters can be seen in Fig. 4.

## 6. Fourier Regularization and Preconditioning

Our learned model weights  $(\{F_c\}, G, B)$  are all periodic  $n \times n$  images. To improve generalization, we want these weights to be small and smooth. In this section we present the general form of the regularization used during training, and we show how this regularization lets us precondition the optimization problem solved during training to find lower-cost minima in fewer iterations. Because this frequency-domain optimization trick applies generally to any optimization problem concerning smooth and periodic images, we will describe it in general terms.

Let us construct an optimization problem with respect to a single  $n \times n$  image  $Z$  consisting of a data term  $f(Z)$  and a regularization  $g(Z)$ :

$$Z^* = \arg \min_Z (f(Z) + g(Z)) \quad (21)$$

We require that the regularization  $g(Z)$  is the weighted sum of squared periodic convolutions of  $Z$  by some filter bank. In our experiments  $g(Z)$  is the weighted sum of the squared difference between adjacent values (similar to a total variation loss [27]) and the sum of squared values:

$$g(Z) = \lambda_1 \sum_{i,j} \left( (Z(i, j) - Z(\text{mod}(i+1, n), j))^2 + (Z(i, j) - Z(i, \text{mod}(j+1, n)))^2 \right) + \lambda_0 \sum_{i,j} Z(i, j)^2 \quad (22)$$

Where  $\lambda_1$  and  $\lambda_0$  are hyperparameters that determine the strength of each smoothness term. We require that  $\lambda_0 > 0$  to prevent divide-by-zero issues during preconditioning.

We use a variant of the standard FFT  $\mathcal{F}_v(\cdot)$  which bijectively maps from some real  $n \times n$  image to a real  $n^2$ -dimensional vector, instead of the complex  $n \times n$  image produced by a standard FFT (See the supplement for a formal description). With this, we can rewrite Eq. 22 as follows:

$$\mathbf{w} = \frac{1}{n} \sqrt{\lambda_1 \left( |\mathcal{F}_v([1, -1])|^2 + |\mathcal{F}_v([1; -1])|^2 \right) + \lambda_0} \\ g(Z) = \mathcal{F}_v(Z)^T \text{diag}(\mathbf{w})^2 \mathcal{F}_v(Z) \quad (23)$$

where the vector  $\mathbf{w}$  is just some fixed function of the definition of  $g(Z)$  and the values of the hyperparameters  $\lambda_1$  and  $\lambda_0$ . The 2-tap difference filters in  $\mathcal{F}_v([1, -1])$  and  $\mathcal{F}_v([1; -1])$  are padded to size  $(n \times n)$  before the FFT. With  $\mathbf{w}$  we can define a mapping between our 2D image space and a rescaled FFT vector space:

$$\mathbf{z} = \mathbf{w} \circ \mathcal{F}_v(Z) \quad (24)$$

Where  $\circ$  is an element-wise product. This mapping lets us rewrite the optimization problem in Eq. 21 as:

$$Z^* = \mathcal{F}_v^{-1} \left( \frac{1}{\mathbf{w}} \left( \arg \min_{\mathbf{z}} \left( f(\mathcal{F}_v^{-1}(\frac{\mathbf{z}}{\mathbf{w}})) + \|\mathbf{z}\|^2 \right) \right) \right) \quad (25)$$

where  $\mathcal{F}_v^{-1}(\cdot)$  is the inverse of  $\mathcal{F}_v(\cdot)$ , and division is element-wise. This reparametrization reduces the complicated regularization of  $Z$  to a simple L2 regularization of  $\mathbf{z}$ , which has a preconditioning effect.

We use this trick during training to reparameterize all model components  $(\{F_c\}, G, B)$  as rescaled FFT vectors, each with their own values for  $\lambda_0$  and  $\lambda_1$ . The effect of this can be seen in Fig. 5, where we show the loss during our two training stages. We compare against naive time-domain

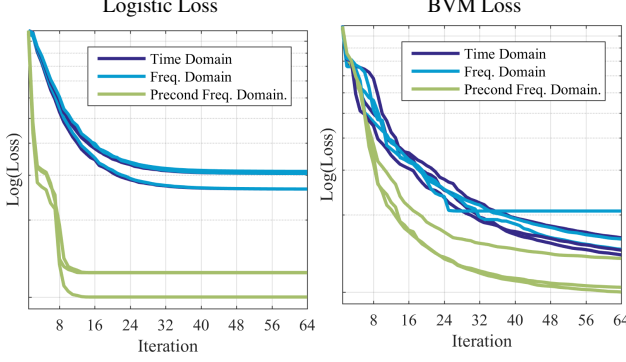


Figure 5: Loss traces for our two stages of training, for three fold cross validation (each line represents a fold) on the Gehler-Shi dataset using LBFGS. Our preconditioned frequency domain optimization produces lower minima at greater rates than are achieved by non-preconditioned optimization in the frequency domain or naive optimization in the time domain.

optimization (Eq. 21) and non-preconditioned frequency-domain optimization (Eq. 25 with  $w = 1$ ). Our preconditioned reformulation exhibits a significant speedup and finds minima with lower losses.

For all experiments, training is as follows: All model parameters are initialized to 0, then we have a convex pre-training step which optimizes Eq. 25 where  $f(\cdot)$  is a logistic loss (described in the supplement) using LBFGS for 16 iterations, and then we optimize Eq. 25 where  $f(\cdot)$  is the non-convex BVM loss in Eq. 18 using LBFGS for 64 iterations.

## 7. Temporal Smoothing

Color constancy algorithms are usually studied in the context of individual images, which are implicitly assumed to be IID. But a practical white balance algorithm must run on a video sequences, and appear to enforce some temporal smoothing of the predicted illuminant to avoid presenting the viewer with an erratically-varying image in the viewfinder. This smoothing must be done intelligently, as an oversmoothed illuminant may appear unresponsive when the environment’s illumination changes rapidly (a colorful light turning on, the camera quickly moving outdoors, etc). Additionally, when faced with multiple valid hypotheses (a blue wall under white light vs a white wall under blue light, etc) we may want to use earlier images to resolve ambiguities. These desiderata of stability, responsiveness, and robustness are all at odds with each other, and so some compromise must be struck.

Our task of constructing a temporally coherent illuminant estimate is aided by the probabilistic nature of the output of our per-frame model, which produces a posterior distribution over illuminants parametrized as a bivariate Gaus-

sian. Let us assume that we have some ongoing estimate of the illuminant and its covariance  $(\mu_t, \Sigma_t)$ . Our model gives us an observed illuminant and covariance  $(\mu_o, \Sigma_o)$ , and with this we update our ongoing estimate by first convolving it with a zero-mean isotropic Gaussian (encoding our prior belief that the illuminant may change over time) and then multiplying that “fuzzed” Gaussian by the observed Gaussian. This results in the following update rule:

$$\Sigma_{t+1} = \left( \left( \Sigma_t + \begin{bmatrix} \alpha & 0 \\ 0 & \alpha \end{bmatrix} \right)^{-1} + \Sigma_o \right)^{-1} \quad (26)$$

$$\mu_{t+1} = \Sigma_{t+1} \left( \left( \Sigma_t + \begin{bmatrix} \alpha & 0 \\ 0 & \alpha \end{bmatrix} \right)^{-1} \mu_t + \Sigma_o \mu_o \right)$$

Where  $\alpha$  is a parameter that defines the expected variance of the illuminant over time. This update resembles a Kalman filter but with a simplified transition model, no control model, and variable observation noise.

This temporal smoothness model is not used in benchmark evaluations, but its effect can be seen in the video included in the supplement.

## 8. Results

We evaluate our technique using two standard color constancy datasets: the Gehler-Shi dataset [18, 28] and on the Cheng et al. dataset [11] (see Tables 1 and 2). For the Gehler-Shi dataset we present several ablations and variants of our model to show the effect of each design decision and to investigate trade-offs between speed and accuracy. Models labeled “full” were run on  $384 \times 256$  16-bit images, while models labeled “thumb” were run on  $48 \times 32$  8-bit images, which are the kind of images that a practical white-balance system embedded on a hardware device might use. Models labeled “4 channel” use the four feature channels used in [4], while models labeled “2 channel” use the two channels we present in Section 5. We also present models which we only use the “pixel channel”  $I$  or the “edge channel”  $E$  as input. All models have a histogram size of  $n = 64$  except for models K and L where  $n$  is varied to show the impact of illuminant aliasing. Two models use “gray world” de-aliasing, and the rest use “gray light” de-aliasing. Gray world de-aliasing appears to be slightly less effective than gray light de-aliasing, except when chroma histograms are heavily aliased, which is why we use it in model K. Model C only has one training stage that minimizes logistic loss for 64 iterations, thereby removing the BVM fitting from training. Model E fixes  $G(i, j) = 1$  and  $B(i, j) = 0$ , thereby removing the model’s ability to reason about the absolute color of the illuminant. Model B was trained only to minimize the data term (ie,  $\lambda_0 = \lambda_1 = 0$  in Eq. 22) while model D uses L2 regularization but not total variation (ie,  $\lambda_1 = 0$  in Eq. 22). For each experiment we

tune all  $\lambda$  hyperparameters to minimize the “average” error during cross-validation, using cyclic coordinate descent.

Model M achieves the lowest-error results results, with a 10% or 13% reduction in error against the previously best-performing published techniques. This improvement in accuracy also comes with a significant speedup compared to previous techniques, with our highest-accuracy model taking 54 milliseconds to process an image, compared to the 520 milliseconds taken by CCC [4] or the 3 seconds (on a GPU) taken by Shi et al. [29]. Model N (our fastest model) has a similar accuracy to [4] and [29] but takes only 1.1 milliseconds to process an image, making it hundreds or thousands of times faster than the current state-of-the art. Additionally, our model appears to be faster to train than the state-of-the-art, though training times for prior work are often not available. All runtimes in Table 1 for our model were computed on an Intel Xeon CPU E5-2680. Runtimes for the “full” model were produced using an unoptimized Matlab implementation, while runtimes for the “thumb” model were produced using a well-optimized Halide [26] implementation running on the CPU (our unoptimized Matlab implementation is 2.37 milliseconds per image).

To demonstrate that our model is a viable automatic white balance system for consumer photography, we ran our Halide code on a 2016 Google Pixel XL using the thumbnail images computed by the device’s camera stack. This implementation ran at 1.44ms per image, which is equivalent to 30 frames per second using  $< 5\%$  of the total compute budget, thereby satisfying our previously-stated speed requirements. A video of our system running in real-time on a phone can be found in the supplement.

## 9. Conclusion

We have presented FFCC, a color constancy algorithm that produces a 10 – 13% reduction in error relative to the best performing prior work, while being 250 – 3000 $\times$  faster than the previous state-of-the-art and fast enough to run in the viewfinder of a consumer camera. In doing this we have introduced the concept of performing convolutional color constancy on a torus, and illuminant aliasing problem this causes, which we solve by also introducing techniques for illuminant de-aliasing and differentiable bivariate von Mises fitting. We have also presented a novel technique for fast Fourier-domain optimization subject to a certain family of regularizers. Unlike past color constancy work, FFCC produces a complete posterior distribution over illuminants, which lets us assess the model’s confidence in its illuminant estimates and also enables a Kalman filter-like temporal smoothing model.

Algorithm	Mean	Med.	Tri.	Best 25%	Worst 25%	Avg.	Test Time	Train Time
Support Vector Regression [16]	8.08	6.73	7.19	3.35	14.89	7.21	-	-
White-Patch [7]	7.55	5.68	6.35	1.45	16.12	5.76	0.16	-
Grey-world [8]	6.36	6.28	6.28	2.33	10.58	5.73	0.15	-
Edge-based Gamut [21]	6.32	5.04	5.43	1.90	13.58	5.40	3.6	1986
1st-order Gray-Edge [30]	5.33	4.52	4.73	1.86	10.03	4.63	1.1	-
2nd-order Gray-Edge [30]	5.13	4.44	4.62	2.11	9.26	4.60	1.3	-
Shades-of-Gray [14]	4.93	4.01	4.23	1.14	10.20	3.96	0.47	-
Bayesian [18]	4.82	3.46	3.88	1.26	10.49	3.86	97	764
General Gray-World [3]	4.66	3.48	3.81	1.00	10.09	3.62	0.91	-
Natural Image Statistics [19]	4.19	3.13	3.45	1.00	9.22	3.34	1.5	10749
CART-based Combination [6]	3.90	2.91	3.21	1.02	8.27	3.14	-	-
Spatio-spectral Statistics [10]	3.59	2.96	3.10	0.95	7.61	2.99	6.9	3159
LSRS [17]	3.31	2.80	2.87	1.14	6.39	2.87	2.6	1345
Intersection-based Gamut [21]	4.20	2.39	2.93	0.51	10.70	2.76	-	-
Pixels-based Gamut [21]	4.20	2.33	2.91	0.50	10.72	2.73	-	-
Bottom-up+Top-down [31]	3.48	2.47	2.61	0.84	8.01	2.73	-	-
Cheng et al. 2014 [11]	3.52	2.14	2.47	0.50	8.74	2.41	0.24	-
Exemplar-based [23]	2.89	2.27	2.42	0.82	5.97	2.39	-	-
Corrected-Moment [13]	2.86	2.04	2.22	0.70	6.34	2.25	0.77	584
Chakrabarti et al. 2015 [9]	2.56	1.67	1.89	0.52	6.07	1.91	0.30	-
Cheng et al. 2015 [12]	2.42	1.65	1.75	0.38	5.87	1.73	0.25	245
CCC [4]	1.95	1.22	1.38	0.35	4.76	1.40	0.52	2168
Shi et al. 2016 [29]	1.90	1.12	1.33	0.31	4.84	1.34	3.0	-
A) FFCC - full, pixel channel only, no illuminant	2.98	1.81	2.00	0.46	7.56	2.06	0.0072	116
B) FFCC - full 2 channels, no regularization	2.36	1.49	1.72	0.56	5.60	1.80	0.028	92
C) FFCC - full 2 channels, no BVM loss	2.16	1.47	1.59	0.75	4.75	1.78	0.027	62
D) FFCC - full 2 channels, no total variation	1.82	1.14	1.26	0.30	4.54	1.29	0.028	79
E) FFCC - full, 2 channels, no illuminant	2.17	1.36	1.51	0.36	5.36	1.54	0.029	81
F) FFCC - full, pixel channel only	2.16	1.29	1.47	0.35	5.48	1.51	0.0074	69
G) FFCC - full, edge channel only	1.99	1.08	1.31	0.32	5.09	1.36	0.023	106
H) FFCC - full, 2 channels, no preconditioning	1.87	1.08	1.26	0.29	4.76	1.29	0.026	81
I) FFCC - full, 2 channels, gray world	1.84	1.06	1.25	0.27	4.67	1.25	0.033	108
J) FFCC - full, 2 channels	1.84	1.05	1.21	0.27	4.70	1.24	0.033	108
K) FFCC - full, 4 channels, $n = 32$ , gray world	2.67	1.32	1.59	0.35	7.31	1.71	0.054	145
L) FFCC - full, 4 channels, $n = 256$	1.78	1.04	1.20	0.27	4.50	1.22	0.060	467
M) FFCC - full, 4 channels	1.78	0.97	1.16	0.27	4.61	1.20	0.054	149
N) FFCC - thumb, 2 channels	1.96	1.21	1.41	0.31	4.84	1.39	0.0011	80

Table 1: Performance on the Gehler-Shi dataset [18, 28]. We present five error metrics and their average (the geometric mean) with the lowest error per metric highlighted in yellow. We present the time (in seconds) for training each model and for evaluating a single image, when available.

Algorithm	Mean	Med.	Tri.	Best 25%	Worst 25%	Avg.
White-Patch [7]	9.91	7.44	8.78	1.44	21.27	7.24
Pixels-based Gamut [21]	5.27	4.26	4.45	1.28	11.16	4.27
Grey-world [8]	4.59	3.46	3.81	1.16	9.85	3.70
Edge-based Gamut [21]	4.40	3.30	3.45	0.99	9.83	3.45
Shades-of-Gray [14]	3.67	2.94	3.03	0.98	7.75	3.01
Natural Image Statistics [19]	3.45	2.88	2.95	0.83	7.18	2.81
Local Surface Reflectance Statistics [17]	3.45	2.51	2.70	0.98	7.32	2.79
2nd-order Gray-Edge [30]	3.36	2.70	2.80	0.89	7.14	2.76
1st-order Gray-Edge [30]	3.35	2.58	2.76	0.79	7.18	2.67
Bayesian [18]	3.50	2.36	2.57	0.78	8.02	2.66
General Gray-World [3]	3.20	2.56	2.68	0.85	6.68	2.63
Spatio-spectral Statistics [10]	3.06	2.58	2.74	0.87	6.17	2.59
Bright-and-dark Colors PCA [11]	2.93	2.33	2.42	0.78	6.13	2.40
Corrected-Moment [13]	2.95	2.05	2.16	0.59	6.89	2.21
Color Dog [2]	2.83	1.77	2.03	0.48	7.04	2.03
Shi et al. 2016 [29]	2.24	1.46	1.68	0.48	6.08	1.74
CCC [4]	2.38	1.48	1.69	0.45	5.85	1.74
Cheng 2015 [12]	2.18	1.48	1.64	0.46	5.03	1.65
M) FFCC - full, 4 channels	1.94	1.35	1.46	0.34	4.51	1.43
N) FFCC - thumb, 2 channels	2.11	1.49	1.62	0.44	4.82	1.61

Table 2: Performance on the dataset from Cheng et al. [11], in the same format as Table 1, excluding runtimes. As was done in [4] we present the average performance (the geometric mean) over all 8 cameras in the dataset.



## References

- [1] E. H. Adelson and A. P. Pentland. Perception as bayesian inference. chapter The Perception of Shading and Reflectance. 1996. 1
- [2] N. Banic and S. Loncaric. Color dog - guiding the global illumination estimation to better accuracy. *VISAPP*, 2015. 8
- [3] K. Barnard, L. Martin, A. Coath, and B. Funt. A comparison of computational color constancy algorithms — part 2: Experiments with image data. *TIP*, 2002. 8
- [4] J. T. Barron. Convolutional color constancy. *ICCV*, 2015. 2, 3, 4, 5, 7, 8
- [5] H. G. Barrow and J. M. Tenenbaum. *Recovering Intrinsic Scene Characteristics from Images*. Academic Press, 1978. 1
- [6] S. Bianco, G. Ciocca, C. Cusano, and R. Schettini. Automatic color constancy algorithm selection and combination. *Pattern Recognition*, 2010. 8
- [7] D. H. Brainard and B. A. Wandell. Analysis of the retinex theory of color vision. *JOSA A*, 1986. 8
- [8] G. Buchsbaum. A spatial processor model for object colour perception. *Journal of the Franklin Institute*, 1980. 8
- [9] A. Chakrabarti. Color constancy by learning to predict chromaticity from luminance. *NIPS*, 2015. 8
- [10] A. Chakrabarti, K. Hirakawa, and T. Zickler. Color constancy with spatio-spectral statistics. *TPAMI*, 2012. 8
- [11] D. Cheng, D. K. Prasad, and M. S. Brown. Illuminant estimation for color constancy: why spatial-domain methods work and the role of the color distribution. *JOSA A*, 2014. 2, 7, 8
- [12] D. Cheng, B. Price, S. Cohen, and M. S. Brown. Effective learning-based illuminant estimation using simple features. *CVPR*, 2015. 8
- [13] G. D. Finlayson. Corrected-moment illuminant estimation. *ICCV*, 2013. 5, 8
- [14] G. D. Finlayson and E. Trezzi. Shades of gray and colour constancy. *Color Imaging Conference*, 2004. 8
- [15] D. H. Foster. Color constancy. *Vision research*, 2011. 1
- [16] B. V. Funt and W. Xiong. Estimating illumination chromaticity via support vector regression. *Color Imaging Conference*, 2004. 8
- [17] S. Gao, W. Han, K. Yang, C. Li, and Y. Li. Efficient color constancy with local surface reflectance statistics. *ECCV*, 2014. 8
- [18] P. Gehler, C. Rother, A. Blake, T. Minka, and T. Sharp. Bayesian color constancy revisited. *CVPR*, 2008. 2, 7, 8
- [19] A. Gijsenij and T. Gevers. Color constancy using natural image statistics and scene semantics. *TPAMI*, 2011. 8
- [20] A. Gijsenij, T. Gevers, and J. van de Weijer. Computational color constancy: Survey and experiments. *TIP*, 2011. 1
- [21] A. Gijsenij, T. Gevers, and J. vande Weijer. Generalized gamut mapping using image derivative structures for color constancy. *IJCV*, 2010. 8
- [22] T. Hamelryck, K. Mardia, and J. Ferkinghoff-Borg. *Bayesian methods in structural bioinformatics*. Springer, 2012. 4
- [23] H. R. V. Joze and M. S. Drew. Exemplar-based color constancy and multiple illumination. *TPAMI*, 2014. 8
- [24] E. H. Land and J. J. McCann. Lightness and retinex theory. *JOSA*, 1971. 1
- [25] K. V. Mardia. Statistics of directional data. *Journal of the Royal Statistical Society, Series B*, 1975. 4
- [26] J. Ragan-Kelley, A. Adams, S. Paris, M. Levoy, S. Amarasinghe, and F. Durand. Decoupling algorithms from schedules for easy optimization of image processing pipelines. *SIGGRAPH*, 2012. 8
- [27] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 1992. 6
- [28] L. Shi and B. Funt. Re-processed version of the gehler color constancy dataset of 568 images. <http://www.cs.sfu.ca/colour/data/>. 2, 7, 8
- [29] W. Shi, C. C. Loy, and X. Tang. Deep specialized network for illuminant estimation. *ECCV*, 2016. 8
- [30] J. van de Weijer, T. Gevers, and A. Gijsenij. Edge-based color constancy. *TIP*, 2007. 8
- [31] J. van de Weijer, C. Schmid, and J. Verbeek. Using high-level visual information for color constancy. *ICCV*, 2007. 8

# Fast Fourier Color Constancy Supplement

Jonathan T. Barron  
barron@google.com

Yun-Ta Tsai  
yuntatsai@google.com

## 1. Pretraining

In the paper we describe the data term for our loss function  $f(\cdot)$  which takes a toroidal PDF  $P(i, j)$ , fits a bivariate von Mises distribution to  $P$ , and then computes the negative log-likelihood of the true white point  $L^*$  under that distribution. This loss is non-convex, and therefore may behave erratically in the earliest training iterations. This issue is compounded by our differentiable BVM fitting procedure, which may be inaccurate when  $P$  has a low concentration, which is often the case in early iterations. For this reason, we train our model in two stages: In the “pretraining” stage we replace the data term in our loss function with a more simple loss: straightforward logistic regression with respect to  $P$  and some ground-truth PDF  $P^*$  (Eq. 2), and then in the second training stage we use the data term described in the paper while using the output of pretraining to initialize the model. Because our regularization is also convex, using this pretraining loss makes our entire optimization problem convex and therefore straightforward to optimize, and (when coupled with our use of LBFGS for optimization instead of some SGD-like approach) also makes training deterministic.

Computing a logistic loss is straightforward: we compute a ground-truth PDF  $P^*$  from the ground-truth illuminant  $L^*$ , and then compute a standard logistic loss.

$$P^*(i, j) = \text{mod} \left( \frac{L_u^* - u_{lo}}{h} - i, n \right) < 1 \quad (1)$$

$$\wedge \text{mod} \left( \frac{L_v^* - v_{lo}}{h} - j, n \right) < 1$$

$$f_{\text{pretrain}}(P) = - \sum_{i,j} P^*(i, j) \log(P(i, j)) \quad (2)$$

This loss behaves very similarly to the loss used in CCC [1], but it has the added benefit of being convex.

## 2. Backpropagation

The bivariate von Mises estimation procedure described in the paper can be thought of as a “layer” in a deep learning architecture, as our end-to-end training procedure requires

that we be able to backpropagate through the fitting procedure and the loss computation. Here we present the gradients of the critical equations described in the paper.

$$\nabla_{\mu} f(\mu, \Sigma) = -2\Sigma^{-1} \left( \begin{bmatrix} L_u^* \\ L_v^* \end{bmatrix} - \mu \right) \quad (3)$$

$$\nabla_{\Sigma} f(\mu, \Sigma) = \Sigma^{-1} - \Sigma^{-1} \left( \begin{bmatrix} L_u^* \\ L_v^* \end{bmatrix} - \mu \right) \left( \begin{bmatrix} L_u^* \\ L_v^* \end{bmatrix} - \mu \right)^T \Sigma^{-1}$$

$$\nabla_{P(i,j)} \mu = \left( \frac{nh}{2\pi} \right) \begin{bmatrix} \frac{x_i \sin(\theta(i)) - y_i \cos(\theta(i))}{x_i^2 + y_i^2} \\ \frac{x_j \sin(\theta(j)) - y_j \cos(\theta(j))}{x_j^2 + y_j^2} \end{bmatrix} \quad (4)$$

$$\nabla_{P(i,j)} \Sigma = h^2 \begin{bmatrix} \bar{i}(\bar{i} - 2E[\bar{i}]), & (\bar{i} - E[\bar{i}])(\bar{j} - E[\bar{j}]) - E[\bar{i}]E[\bar{j}] \\ (\bar{i} - E[\bar{i}])(\bar{j} - E[\bar{j}]) - E[\bar{i}]E[\bar{j}], & \bar{j}(\bar{j} - 2E[\bar{j}]) \end{bmatrix}$$

By chaining these gradients together we can backpropagate the gradient of the loss back onto the input PDF  $P$ . Backpropagating through the softmax operation and the convolution (and illumination gain/bias) is straightforward and so is not detailed here.

## 3. Results

Because our model produces a complete posterior distribution over illuminants in the form of a covariance matrix  $\Sigma$ , each of our illuminant estimates comes with a measure of confidence in the form of the entropy:  $\frac{1}{2} \log |\Sigma|$  (ignoring constant factors). A low entropy suggests a tight concentration of the output distribution, which tends to be well-correlated with a low error. To demonstrate this we present a novel error metric, which is twice the area under the curve formed by ordering all images (the union of all test-set images from each cross-validation fold) by ascending entropy, taking the cumulative sum of the sorted error, and normalizing by the number of images. In Figure 1 we visualize this error metric and show that our entropy-ordered error is substantially lower than the mean error for both of our datasets, which shows that a low entropy is suggestive of a low error. We are not aware of any other color constancy technique which explicitly predicts a confidence measure, and so we do not compare against any existing technique, but it can be demonstrated that if the entropy used to sort

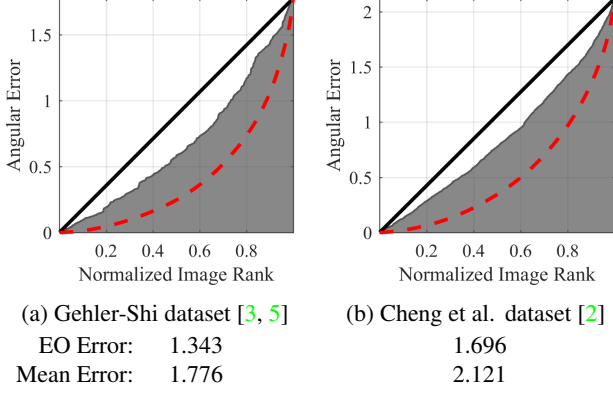


Figure 1: By sorting each image by the entropy of its posterior distribution we can show that entropy correlates with error. Here we sort the images by ascending entropy and plot the cumulative sum of the error, filling in the area under that curve with gray. If entropy was not correlated with error we would expect the area under the curve to match the black line, and if entropy was perfectly correlated with error then the area under the curve would exactly match the dashed red line. We report twice the area under the curve as “entropy-ordered” error (mean error happens to be twice the area under the diagonal line).

error is decorrelated with error (or, equivalently, if the error cannot be sorted due to the lack of the means to sort it) that entropy-ordered error will on average be equal to mean error.

To allow for a better understanding of our model’s performance, we present images from the Gehler-Shi dataset [3, 5] (Figures 2-11) and the Canon 1Ds MkIII camera from the Cheng et al. dataset [2] (Figures 12-16). For each dataset we perform three-fold cross validation, and with that we produce output predictions for each image along with an error measure (angular RGB error) and an entropy measure (the entropy of the covariance matrix of our predicted posterior distribution over illuminants). The images chosen here were selected by sorted images from each dataset by increasing error and evenly sampling images according to that ordering (10 from Gehler-Shi, 5 from the smaller Cheng dataset). This means that the first image in each sequence is the lowest error image, and the last is the highest. The rendered images include the color checker used in creating the ground-truth illuminants used during training, but it should be noted that these color checkers are masked out when these images are used during training and evaluation. For each image we present: a) the input image, b) the predicted bivariate von Mises distribution over illuminants, c) our estimated illuminant and white-balanced image (produced by dividing the estimated illuminant into the input image), and d) the ground-truth illuminant and white-balanced image.

Our log-chroma histograms are visualized using gray light de-aliasing to render each  $(u, v)$  coordinate’s color, with a blue dot indicating the location/color of the ground-truth illuminant, a red dot indicating our predicted illuminant  $\mu$  and a red ellipse indicating the predicted covariance of the illuminant  $\Sigma$ . The bright lines in the histogram indicate the locations where  $u = 0$  or  $v = 0$ . The reported entropy of the covariance  $\Sigma$  corresponds to the spread of the covariance (low entropy = small spread). We see that our low error predictions tend to have lower entropies, and vice versa, confirming our analysis in Figure 1. We also see that the ground-truth illuminant tends to lie within the estimated covariance matrix, though not always for the largest-error images.

In Figure 17 we visualize a set of images taken from a Nexus 6 in the HDR+ mode [4] after being white-balanced by the “thumb” version of our algorithm (Model M in the main paper).

#### 4. Real Bijective FFT

In the paper we describe  $\mathcal{F}_v(Z)$ , a FFT variant that takes the 2D FFT of a  $n \times n$  real-valued 2D image  $Z$  and then linearizes it into a real-valued vector with no redundant values. Having this FFT-like one-to-one mapping between real 2D images and real 1D vectors enables our frequency-domain preconditioner.

Our modified FFT function is defined as:

$$\mathcal{F}_v(Z) = \begin{bmatrix} \text{Re}(\mathcal{F}(Z)(0 : n/2, 0)) \\ \text{Re}(\mathcal{F}(Z)(0 : n/2, n/2)) \\ \text{Re}(\mathcal{F}(Z)(0 : (n-1), 1 : (n/2-1))) \\ \text{Im}(\mathcal{F}(Z)(1 : (n/2-1), 0)) \\ \text{Im}(\mathcal{F}(Z)(1 : (n/2-1), n/2-1)) \\ \text{Im}(\mathcal{F}(Z)(0 : (n-1), 1 : (n/2-1))) \end{bmatrix} \quad (5)$$

Where  $\mathcal{F}(Z)(i, j)$  is the complex number at the zero-indexed  $(i, j)$  position in the FFT of  $Z$ ,  $\text{Re}(\cdot)$  and  $\text{Im}(\cdot)$  extract real and imaginary components, respectively, and range indexing is inclusive. The output of  $\mathcal{F}_v(Z)$  is an  $n^2$ -dimensional vector, as it must be for our mapping to necessarily preserve all of the FFT coefficients with no redundancy. To preserve the scale of the FFT through this mapping, we scale  $\mathcal{F}_v(Z)$  by  $\sqrt{2}$ , ignoring the entries that correspond to:

$$\begin{matrix} \text{Re}(\mathcal{F}(Z)(0, 0)) & \text{Re}(\mathcal{F}(Z)(0, n/2)) \\ \text{Re}(\mathcal{F}(Z)(n/2, 0)) & \text{Re}(\mathcal{F}(Z)(n/2, 0)) \end{matrix} \quad (6)$$

This scaling ensure that the magnitude of  $Z$  is preserved:

$$\|\mathcal{F}_v(Z)\|^2 = |\mathcal{F}(Z)|^2 \quad (7)$$

To compute the inverse of  $\mathcal{F}_v(\cdot)$  we undo this scaling, undo the vectorization by filling in a subset of the elements of

$\mathcal{F}(Z)$  from the vector representation, set the other elements of  $\mathcal{F}(Z)$  such that Hermitian symmetry holds, and the invert the FFT.

## 5. Color Rendering

All images (and all log-chroma histogram visualizations) are rendered by applying the RGB gains implied by the estimated illuminant, applying some color correction matrix (CCM) and then applying an sRGB gamma-correction function (the  $C_{\text{linear}}$  to  $C_{\text{srgb}}$  mapping in <http://en.wikipedia.org/wiki/SRGB>. For each camera in the datasets we estimated our own CCMs using the imagery, which we present here. These CCMs do not affect our illuminant estimation or our results, and are only relevant to our visualizations. Each CCM is estimated through an iterative least-squares process in which we alternately: 1) estimate the ground-truth RGB gains for each image from a camera by solving a least-squares system using our current CCM estimate, and 2) use the estimated gains to estimate a row-normalized CCM using a constrained least-squares solve. Our estimated ground-truth gains are not used in this paper. For the ground-truth sRGB colors of the Macbeth color chart we use the hex values provided here: <http://en.wikipedia.org/wiki/ColorChecker#Colors> which we convert to linear RGB before calibration.

$$\begin{aligned} \text{Cheng, SamsungNX2000} & \begin{bmatrix} 1.5770 & -0.4351 & -0.1419 \\ -0.1747 & 1.5225 & -0.3477 \\ 0.0573 & -0.6397 & 1.5825 \end{bmatrix} \\ \text{Cheng, SonyA57} & \begin{bmatrix} 1.5963 & -0.5545 & -0.0418 \\ -0.1343 & 1.5331 & -0.3988 \\ 0.0563 & -0.4026 & 1.3463 \end{bmatrix} \end{aligned}$$

## References

- [1] J. T. Barron. Convolutional color constancy. *ICCV*, 2015. 1
- [2] D. Cheng, D. K. Prasad, and M. S. Brown. Illuminant estimation for color constancy: why spatial-domain methods work and the role of the color distribution. *JOSA A*, 2014. 2
- [3] P. Gehler, C. Rother, A. Blake, T. Minka, and T. Sharp. Bayesian color constancy revisited. *CVPR*, 2008. 2
- [4] S. W. Hasinoff, D. Sharlet, R. Geiss, A. Adams, J. T. Barron, F. Kainz, J. Chen, and M. Levoy. Burst photography for high dynamic range and low-light imaging on mobile cameras. *SIGGRAPH Asia*, 2016. 2, 7
- [5] L. Shi and B. Funt. Re-processed version of the gehler color constancy dataset of 568 images. <http://www.cs.sfu.ca/colour/data/>. 2

$$\begin{aligned} \text{GehlerShi, Canon1D} & \begin{bmatrix} 2.2310 & -1.5926 & 0.3616 \\ -0.1494 & 1.4544 & -0.3050 \\ 0.1641 & -0.6588 & 1.4947 \end{bmatrix} \\ \text{GehlerShi, Canon5D} & \begin{bmatrix} 1.7494 & -0.8470 & 0.0976 \\ -0.1565 & 1.4380 & -0.2815 \\ 0.0786 & -0.5070 & 1.4284 \end{bmatrix} \\ \text{Cheng, Canon1DsMkIII} & \begin{bmatrix} 1.7247 & -0.7791 & 0.0544 \\ -0.1436 & 1.4632 & -0.3195 \\ 0.0589 & -0.4625 & 1.4037 \end{bmatrix} \\ \text{Cheng, Canon600D} & \begin{bmatrix} 1.8988 & -0.9897 & 0.0909 \\ -0.2058 & 1.6396 & -0.4338 \\ 0.0749 & -0.7030 & 1.6281 \end{bmatrix} \\ \text{Cheng, FujifilmXM1} & \begin{bmatrix} 1.4183 & -0.2497 & -0.1686 \\ -0.2230 & 1.6449 & -0.4219 \\ 0.0785 & -0.5980 & 1.5195 \end{bmatrix} \\ \text{Cheng, NikonD5200} & \begin{bmatrix} 1.3792 & -0.3134 & -0.0659 \\ -0.0826 & 1.3759 & -0.2932 \\ 0.0483 & -0.4553 & 1.4070 \end{bmatrix} \\ \text{Cheng, OlympusEPL6} & \begin{bmatrix} 1.6565 & -0.4971 & -0.1595 \\ -0.3335 & 1.7772 & -0.4437 \\ 0.0895 & -0.7023 & 1.6128 \end{bmatrix} \\ \text{Cheng, PanasonicGX1} & \begin{bmatrix} 1.5629 & -0.5117 & -0.0512 \\ -0.2472 & 1.7590 & -0.5117 \\ 0.1395 & -0.8945 & 1.7550 \end{bmatrix} \end{aligned}$$



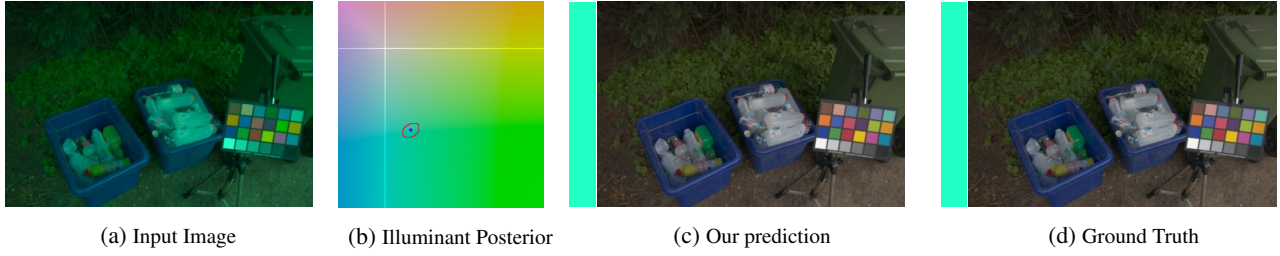


Figure 2: A result from the Gehler-Shi dataset. Error =  $0.02^\circ$ , entropy =  $-6.48$

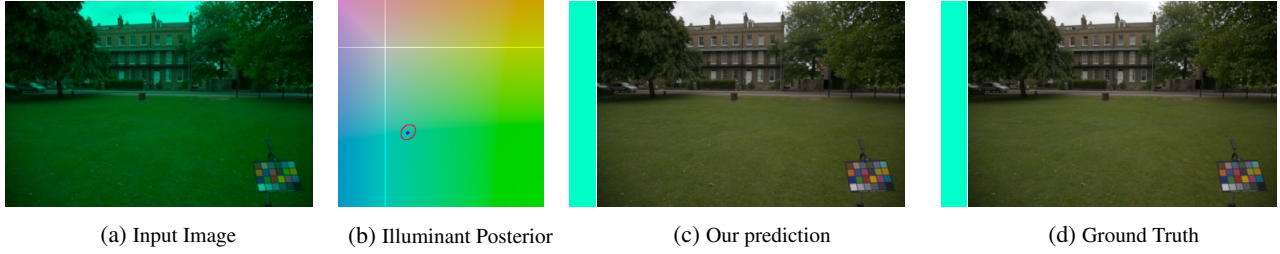


Figure 3: A result from the Gehler-Shi dataset. Error =  $0.26^\circ$ , entropy =  $-6.55$

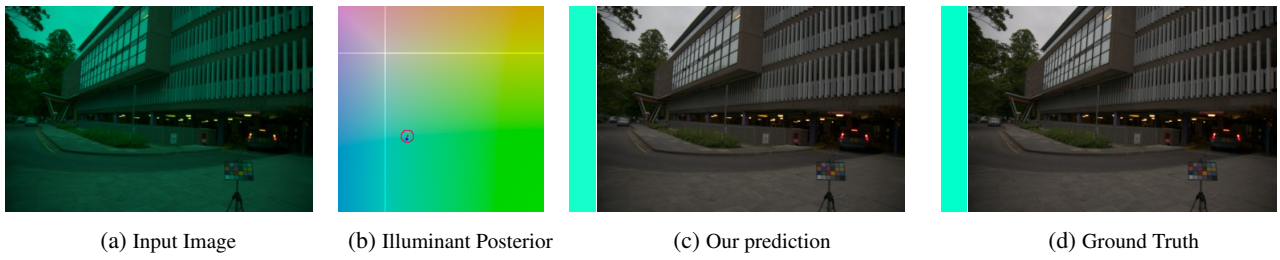


Figure 4: A result from the Gehler-Shi dataset. Error =  $0.46^\circ$ , entropy =  $-6.91$

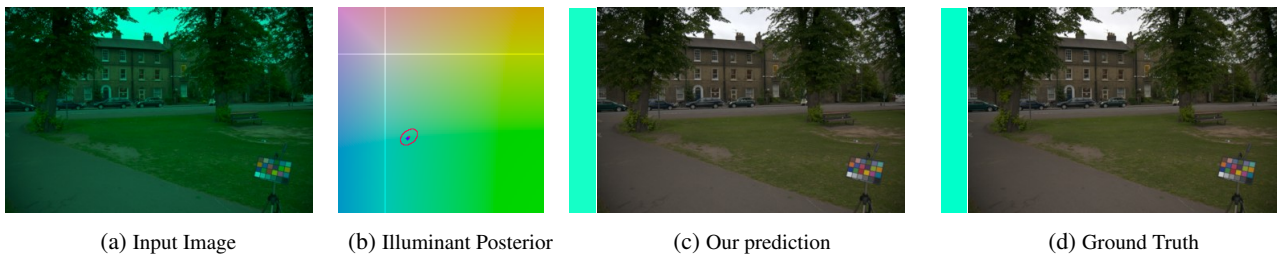


Figure 5: A result from the Gehler-Shi dataset. Error =  $0.63^\circ$ , entropy =  $-6.37$

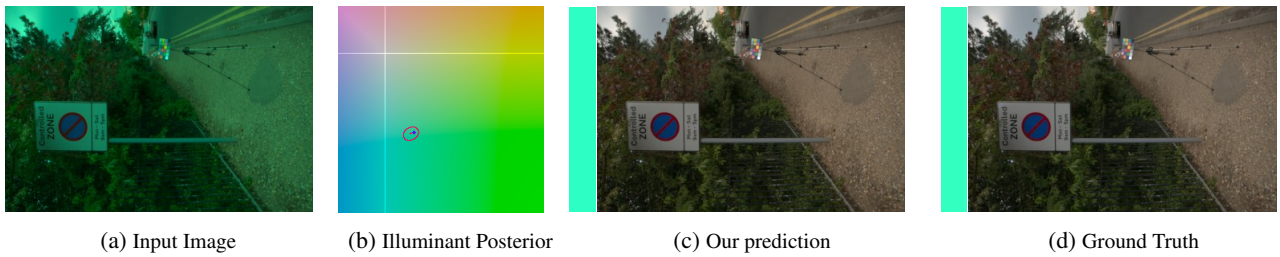


Figure 6: A result from the Gehler-Shi dataset. Error =  $0.83^\circ$ , entropy =  $-6.62$

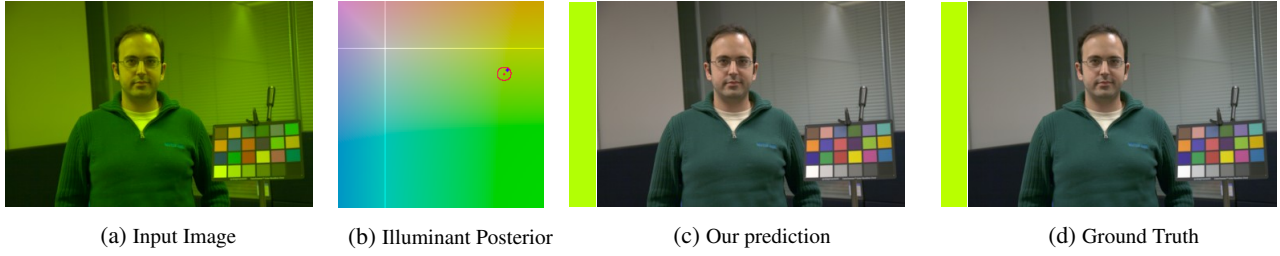


Figure 7: A result from the Gehler-Shi dataset. Error =  $1.19^\circ$ , entropy =  $-6.71$

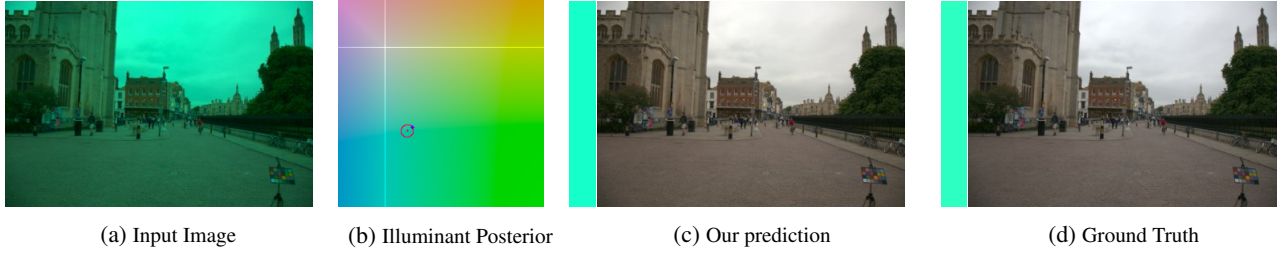


Figure 8: A result from the Gehler-Shi dataset. Error =  $1.61^\circ$ , entropy =  $-6.88$

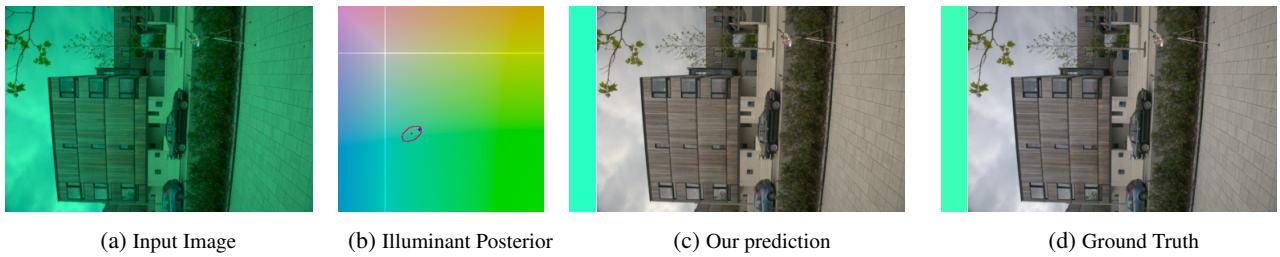


Figure 9: A result from the Gehler-Shi dataset. Error =  $2.35^\circ$ , entropy =  $-6.32$

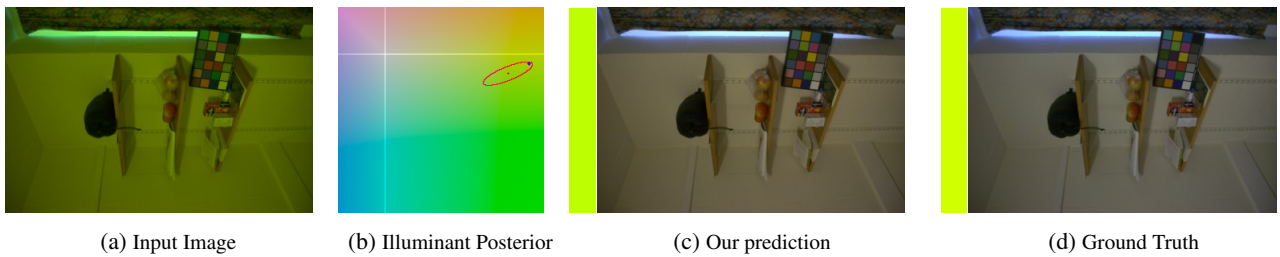


Figure 10: A result from the Gehler-Shi dataset. Error =  $3.84^\circ$ , entropy =  $-5.28$

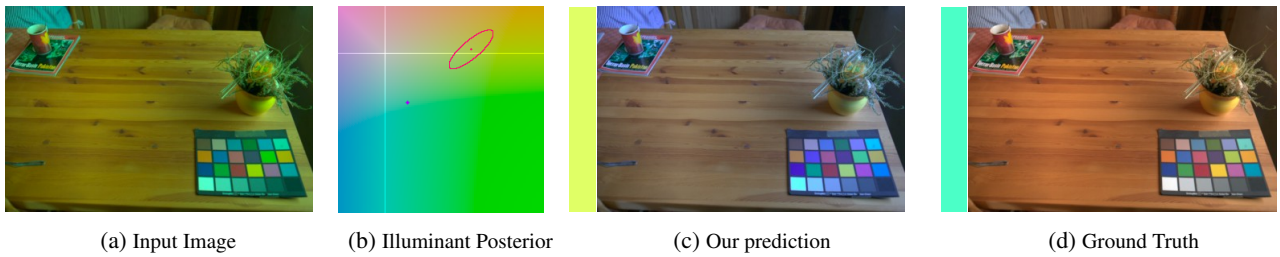
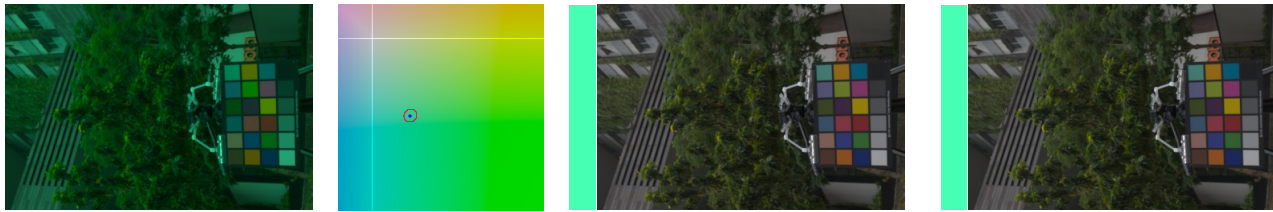


Figure 11: A result from the Gehler-Shi dataset. Error =  $21.64^\circ$ , entropy =  $-4.95$





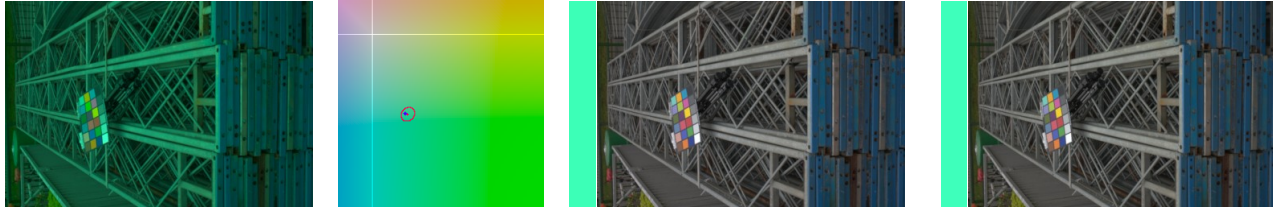
(a) Input Image

(b) Illuminant Posterior

(c) Our prediction

(d) Ground Truth

Figure 12: A result from the Cheng dataset. Error =  $0.12^\circ$ , entropy =  $-6.82$



(a) Input Image

(b) Illuminant Posterior

(c) Our prediction

(d) Ground Truth

Figure 13: A result from the Cheng dataset. Error =  $0.64^\circ$ , entropy =  $-6.69$



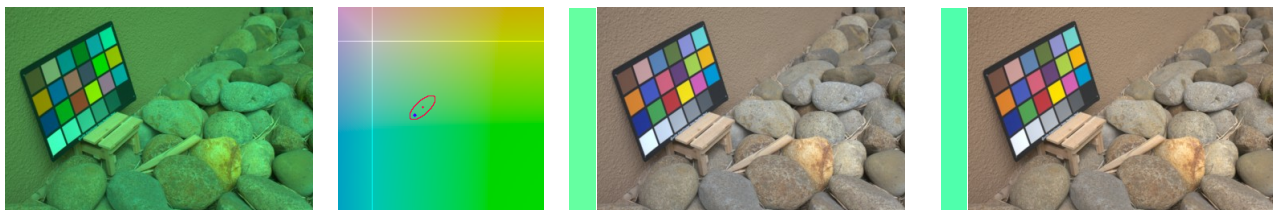
(a) Input Image

(b) Illuminant Posterior

(c) Our prediction

(d) Ground Truth

Figure 14: A result from the Cheng dataset. Error =  $1.37^\circ$ , entropy =  $-6.48$



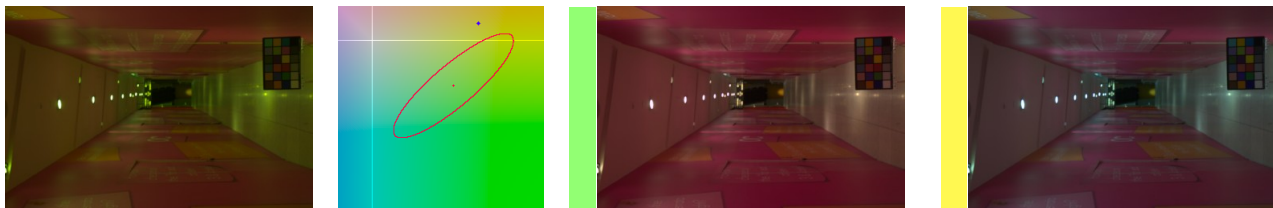
(a) Input Image

(b) Illuminant Posterior

(c) Our prediction

(d) Ground Truth

Figure 15: A result from the Cheng dataset. Error =  $2.69^\circ$ , entropy =  $-5.82$



(a) Input Image

(b) Illuminant Posterior

(c) Our prediction

(d) Ground Truth

Figure 16: A result from the Cheng dataset. Error =  $17.85^\circ$ , entropy =  $-3.04$



Figure 17: A sampling of unedited HDR+[4] images from a Nexus 6, after being processed with the “thumb” version of our algorithm.